(4)

NUSC Technical Report 6760
1 February 1989

Information Gathering for Adaptable Decision-Making

S. S. Kirschenbaum
Combat Control Systems Department

Naval Underwater Systems Center
Newport, Rhode Island / New London, Connecticut

PREFACE

The research for the Adaptable Expert Systems project was supported in part by the Chief of Naval Research, program manager J.G. Smith (OCNR-121). The project was designed to investigate and demonstrate the feasibility of an expert system that adapts automatically to the needs of its user.

The author gratefully acknowledges the contributions of L.M. Cabral, A.H. Silva, F.J. O'Brien and, especially, S.E. Bill-Wray (all of Code 2212). The comments and research assistance of C.C. Celebucki (Code 2212) are also greatly appreciated.

REVIEWED AND APPROVED: 1 FEBRUARY 1989

J.R. Short
Head, Combat Control Systems Department

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) TR 6760	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION Naval Underwater Systems Ctr 6b. OFFICE SYMBOL (If applicable) Code 2212	7a. NAME OF MONITORING ORGANIZATION
6c. ADDRESS (City, State, and ZIP Code) Newport Laboratory Newport, RI 02841	7b. ADDRESS (City, State, and ZIP Code)
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Chief of Naval Research 8b. OFFICE SYMBOL (If applicable) OCNR-121	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
8c. ADDRESS (City, State, and ZIP Code) Washington, DC 20362	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. / PROJECT NO. / TASK NO / WORK UNIT ACCESSION NO

11. TITLE (Include Security Classification)
INFORMATION GATHERING FOR ADAPTABLE DECISION-MAKING

12. PERSONAL AUTHOR(S)
Kirschenbaum, S.S.

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final	FROM ____ TO ____	89-02-01	35

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	
15	01		Anti-Submarine Warfare Adaptable Expert System Decision-Making

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

The Adaptable Expert Systems project was designed to investigate and demonstrate the feasibility of an expert system that adapts automatically to the needs of its user. As part of this project, there was an interdisciplinary collaboration between experimental psychology and computer science. The psychological research is documented in this report and includes behavioral observations and a controlled experimental investigation of information-gathering strategies for command decision-making.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a. NAME OF RESPONSIBLE INDIVIDUAL S.S. Kirschenbaum	22b. TELEPHONE (Include Area Code) (401) 841-2648 22c. OFFICE SYMBOL Code 2212

DD FORM 1473, 84 MAR
83 APR edition may be used until exhausted
All other editions are obsolete

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

INFORMATION GATHERING FOR ADAPTABLE DECISION-MAKING

INTRODUCTION

The design of command decision support aids requires an understanding of the human decision process and the circumstances under which command decisions are made. For such aids to represent an advancement over existing technology and to meet with user acceptance, they must be robust, reliable, flexible, and do what the user wishes done. This is a difficult task, made more difficult by the variety of possible users and circumstances. One solution to this problem is the creation of an aid that adapts to the needs of its user. The research reported here is part of the Adaptable Expert Systems (AES) project, which was designed to investigate the feasibility of such a decision aid.

ADAPTABILITY: HUMAN AND COMPUTER

Adaptability is a new capability for an expert system, but a common feature of human behavior. People naturally and automatically respond differently to a child than to an adult, to driving a truck than to driving a sports car, to being a student than to being an instructor, to an angry coworker than to a happy one. These differences are adaptations to cues received from the situations and people involved in the interaction. For example, Isaacs and Clark[1] have shown that, in response to cues from their partners, experts and novices adjust their conversational references and perspectives to the perceived level of expertise of their partner. Moreover, they take different roles, with experts supplying information and novices acquiring it over time.

An expert system that responds adaptively is built along the same lines; that is, it responds to cues from the environment or the user. The computer, like the human, could be made to adapt to differences in situation, technology, or people (users). The payoff of human-like adaptability in a decision aid is expected to be in two areas. First, as with the human, adaptability should make the response of the system more appropriate for the user and situation and, therefore, contribute to better user-system performance. Secondly, appropriate responses from the system should help alleviate the problem of user acceptance commonly found with expert systems. The ultimate responsibility for any decision always remains with the human, and the consequences of a good or poor decision will fall on human shoulders. Thus, any decision aid must prove its value before it will be accepted as helpful. This is even more true of an expert system that gives advice than it is of an aid that only helps with information management and algorithm calculation. The reasons for this problem include lack of user control, the hidden nature of the reasoning, and the rigidity of the solution and its presentation to the user. Adaptability is expected to help win user acceptance by using an interactive and collaborative process and by adapting reasoning and presentation to user needs. An expert user may require information management and confirmation/disconfirmation of his own decision. A less experienced decision-maker may need explanation and guidance toward a correct decision. Thus, the adaptable expert system, like the human partner in decision-making,[1] should adjust its role and output to the experience level of its user if its advice is to be accepted.

The range of things to which a system might adapt is very wide. Within the submarine domain alone, it includes environmental conditions, threat profile, rules of engagement, own ship resources, and both between-user and within-user differences. For the purpose of this project, adapting to between-user differences resulting from the level of experience was chosen as a major focus. This is a relevant and meaningful choice because of differences in officer training and experience and the need to support many levels of command decision-maker stand-ins, as well as the commanding officers (CO) themselves. The officer performing as the approach officer (AO), CO or not, may have a speciality in any one of several areas, varying from sonar to engineering. His previous training and experience may have prepared him to overvalue some sources of information or not have prepared him to appreciate fully the value of others in certain situations. Moreover, the situation itself may be new to him. Experience in the Mediterranean may not always generalize to the North Atlantic.

The approach taken on this project was an interdisciplinary collaboration between experimental psychology and computer science. The role of the psychologist was to develop an understanding of the effects of differences in levels of experience on the needs and information search behavior of users. The role of the computer scientist was to develop a strategy whereby an adaptable expert system should recognize user actions and respond appropriately. The psychological research is reported herein, and includes a brief review of the relevant literature and reports of both behavioral observations and a controlled experimental investigation of information-gathering strategies for command decision-making. The results of this research have been applied to the adaptable expert system, named Ranger, which exists in the laboratory. The computer science portion of the AES project has been documented by Cary.[2]

PSYCHOLOGICAL LITERATURE

This research brings together two related themes in the recent literature. The first of these is the work on decision-making strategies done by Payne;[3] Klayman;[4,5] Einhorn and Hogarth;[6,7,8] Einhorn, Kleinmuntz, and Kleinmuntz;[9] and Svenson.[10] The second theme is the work on novice-expert differences in problem-solving, which goes back to deGroot.[11] As Newell[12] notes, decision-making is a form of problem-solving. Good decision-making, like any other problem-solving activity, is an acquired skill. Expert decision-makers are rewarded with success in many arenas (e.g., financial, social, political, and military). As with the problem-solving research,[13,14] the differences between novice and experienced decision-makers offer tantalizing glimpses into the details of the decision-making process and the significance of hypothesized variables. They may also provide suggestions for adaptive decision aiding and improvements in training.

This research was conducted with a submarine problem: deciding how to respond to a passive sonar target. In this problem, the submarine operator also must gather information and decide how to respond to other objects. The decision-making involved in such a scenario is not the slow, considered, economic decision-making of classic normative[15,16] or descriptive[17] decision

theory. In the class of decisions modeled by classic theories, the situation is largely known (or knowable) before the decision is made. In the dynamic case, such as the decision situation of the submariner deciding how to respond to a sonar contact, situational information is ambiguous and rapidly changing.

The process of decision-making has been discussed in Kirschenbaum.[18] Based on the SHOR paradigm,[19] the SHORE (stimulus, hypothesis, option, response, evaluation) model was developed to be descriptive of command decision-making in the ambiguous, dynamic submarine context. This process begins when some change occurs in the situation or decision-maker needs and goals. The decision-maker must assess the available situational information (stimulus) to determine the current state of the world (hypothesis). This might require repeated information-gathering actions and might be difficult if the available information is ambiguous, incomplete, or missing. Either in parallel or after reasonable satisfaction with the situation assessment, response options are generated (option). These are evaluated against goals and expected outcomes. After the decision response is made (response), the situational consequences are observed. This feedback (evaluation) is added to the individual's knowledge base and used to confirm or modify beliefs about the current situation.

Frequently, situation assessment is critical to decision-making because the situation itself either determines or severely limits the response options. For example, Kahneman and Tversky[17,20] and Tversky and Kahneman[21,22] have found that elements of how the situation is presented (and, by extension, assessed by the decision-maker) strongly influence the decision response. For example, in an automobile, the location, direction of movement, and speed of other objects (cars, pedestrians, and other obstacles) can fully determine the driver's response.

The situation assessment process can be further decomposed into selective attention to stimuli, information search, matching of information with previously learned patterns of events, and further search to narrow and/or confirm the selection of a match. The information-gathering process is of special interest to the current research for several reasons. First, it can be a critical step in the decision-making process because of the connections among available information, experience, and good decision-making. Even the best person can do no better than chance without information. Second, information search is one of only two opportunities to observe the steps of decision-making. Moreover, it is a very useful way of learning about the thought process because it occurs many times over the course of the process. The decision response itself occurs only once, at the end of the process. Third, decision support systems, such as the proposed AES, are primarily information analysis and display systems. The user gathers the presented information, including suggested actions and reasons/explanations, and makes the actual decision. Thus, knowing how different classes of decision-makers gather information can lead to better information presentation.

The situation assessment process is analogous to problem classification and assessment in problem-solving tasks. As with expert problem-solving,[23,24] the skilled decision-maker must compare the information gathered in the new instance to patterns of information, schema,[25] productions,[13] or mental

models[26] found in known rules and previously encountered situations. He or she can then act from a knowledge of the previous response(s) and outcome(s), taking into account any necessary alterations in action due to differences in the situation and/or current goal state. If there is a good match and if the previous encounter ended in a way that is satisfactory to the decision-maker's current state, then the selection of response is relatively easy. If only a partial match is found, the decision-maker may infer that the new situation is an instance of a known class of situations and can construct a new mental model of the situation, incorporating variations from the previously encountered version(s). If either the current situation or state of the organism is radically different, then the generation and the selection of the decision response is less predictable.

LEVEL OF EXPERIENCE DIFFERENCES

The experienced decision-maker, like the expert problem-solver,[25,27] looks for specific types of information during the initial situation assessment phase. The novice does not have such a set of information-seeking strategies in place and may neglect to gather all the relevant information or may waste time on less relevant items. In the limited times availables, he or she simply cannot decide what is most important. Furthermore, for an experienced person, the response is drawn from a very limited set and is a variant of one that worked in previous, similar situations. When the proper script, mental model, or expectancy set is identified, one can apply the rule "standard action for standard situation."[28] In contrast, for the novice, the current elements do not fit any known situation, but may suggest several possible situations. Because it is not bounded by the experience of a successful or unsuccessful outcome, this is a larger set than would be considered by a more experienced person. This might explain why researchers have found that experts not only solve problems more quickly, but consider fewer alternative solution paths than do novices.[23,25]

Extensive research into differences between novices and experts has shown interesting differences in problem solving in such diverse domains as chess,[29] physics,[14,30] activity,[13] and computer programming.[31] To summarize this research, expert behavior is more systematic and efficient in all domains. Experts and novices categorize and represent problems differently, and this contributes to differences in solution behavior. Experts appear to consider fewer solution options and to make use of larger chunks of solution behavior. The expert representation of the problem space appears to include both physical objects and abstract "forces" and similar domain constructs while novices make use of only physical objects.

Chase and Simon[29] estimate that an expert takes 20,000 hours (about 10 years) to master a domain. By this definition, there are very few experts in any domain. Some submarine COs are among this select group. The training program for submarine command includes years of experience at sea supplemented by years of education and training in classrooms and simulators. Many COs have advanced degrees in various aspects of ocean engineering and have spent time as instructors of submarine operations. They are thoroughly familiar with relevant naval doctrine and have the experience to recognize and respond appropriately to variations and exceptions of the standard situations covered by doctrine.

The so-called novice/expert difference is really not a dichotomy, but a continuum. One intermediate point along this continuum may be represented by the rule-bound decision-maker, i.e., one who has learned the rules for a particular task, but has not had extensive experience applying them. One characteristic of rules is that they work best in ideal, error-free conditions. Thus, the rule-bound decision-maker should be able to perform well in a task that conforms to the rules, but lacks the flexible expectancy set that develops with feedback from realistic conditions over the course of many exposures. One example of this group is composed of those instructors (e.g., graduate student teaching assistants or instructors at the U.S. Navy Submarine School) who teach, but have not had the extensive experience in their chosen field that characterizes an expert. The performance differences between such instructors and experienced decision-makers can help define the contributions of rule learning and actual experience. Instructors become very proficient at enunciating (and demonstrating) the rules of a task, but may lack the experience to apply them in nonstandard situations.

In contrast to both experts and instructors, beginning students, as in all fields, are just acquiring a familiarity with the vocabulary and concepts necessary to understand the standard rules of operation. They have not yet acquired the skills and knowledge to operate confidently in this new domain.

BEHAVIORAL OBSERVATIONS

The first phase in the experimental process was to gain a familiarity with the activities and behavior of the crew in a submarine command center. The purposes were to learn more about the command decision process and to generate hypotheses about the effect of experience level on domain-specific information-gathering behavior. This process took two forms: participation in a course for Mk 117 system operators and behavioral observations in tactical simulators.

Participation in the course provided the experimenters with a basic familiarity with the sources of information available for command decision-making. It also highlighted some of the difficulties in sources of ambiguity. Lastly, because of the composition of the student body, this course provided an opportunity to informally observe both novice and experienced naval personnel at work.

Observations indicated that operation of the system depends upon mastery of information about the system, manual skill, and ability to visualize the positions of own ship and target. Those students with at-sea experience appeared to perform better at this last component during laboratory exercises, although they did not necessarily score better on tests of information or have better motor skills.

Note that the focus of these classes was on operation of the system, not on command decision-making. To better understand the command process and develop hypotheses, discussions were held with experts on submarine command decision-making. In addition, formal behavioral observations were conducted in tactical trainers. These trainers are simulators of actual attack submarine

command centers in which controlled exercise scenarios can be run and replayed for analysis. During these observations, existing crews were training at their regular battle station assignments. Observations took place over the course of 10 scenarios. Both the principle decision-maker (the CO) and his senior officers were observed during information gathering and decision-making. Behavior was recorded on a standard behavioral observation form.

Results, summarized for both the CO and senior officers, were used to develop the search strategy hypotheses discussed herein. In brief, COs made use of a variety of information from all sources, although there were differences among COs observed. There was a great deal of computer-processed information available; however, COs appeared to weight heavily such relatively unprocessed data as sonar traces and paper plots as confirmation of computer-generated solutions. In contrast, other officers, who were assigned to specific areas, did not appear to do such compensatory weighting. They seemed to weight their advice toward only their own area of responsibility. Thus, a weapons officer considered firing position and the fire control coordinator, who is responsible for pinpointing the location of a target, recommended a maneuver that would optimize the ship's position for target motion analysis. These observations led to the predictions discussed in the following section.

INFORMATION SEARCH STRATEGIES EXPERIMENT

Several information search strategies have been reported for the choice decision situation.[4,5,10] It has been shown by Klayman[4,5] that humans efficiently adjust their search strategies for such context effects as time available, quantity of information, and task. One can propose related information-seeking strategies in the dynamic decision situation. Furthermore, one can predict that the efficiency of these search strategies will vary with the experience level of the decision-maker. A strategy that is efficient for the expert may not be efficient for the novice.

In the current work, seven decision strategies were defined and predictions made as to their use by experienced or novice submarine decision-makers. These strategies, while specific to the submarine decision problem, are representative of more general information search strategies. They should provide an indication of differences due to a level of experience that can be generalized to other dynamic, ambiguous decision problems.

The first three of these strategies were defined by frequency of access, the second two by single-stage transitions between two data items, and the last two by patterns composed of one or more transitions. For clarity, the term "look" is used to denote a single inspection of any item of information, such as "computed bearing" or "tracker."

The experiment used a process-tracing approach to test predictions about the use of these strategies by three classes of decision-makers: (1) those with both experience and knowledge of rules, (2) rule-bound instructors with only limited experience in applying the rules, and (3) novice, student decision-makers. In the experiment, information typically found in a

submarine combat control center was presented in a matrix of item-by-time-slice data (figure 1). The subjects used a mouse to select desired information. In the following sections, predictions for student and CO groups are discussed, as well as the variations for instructors.

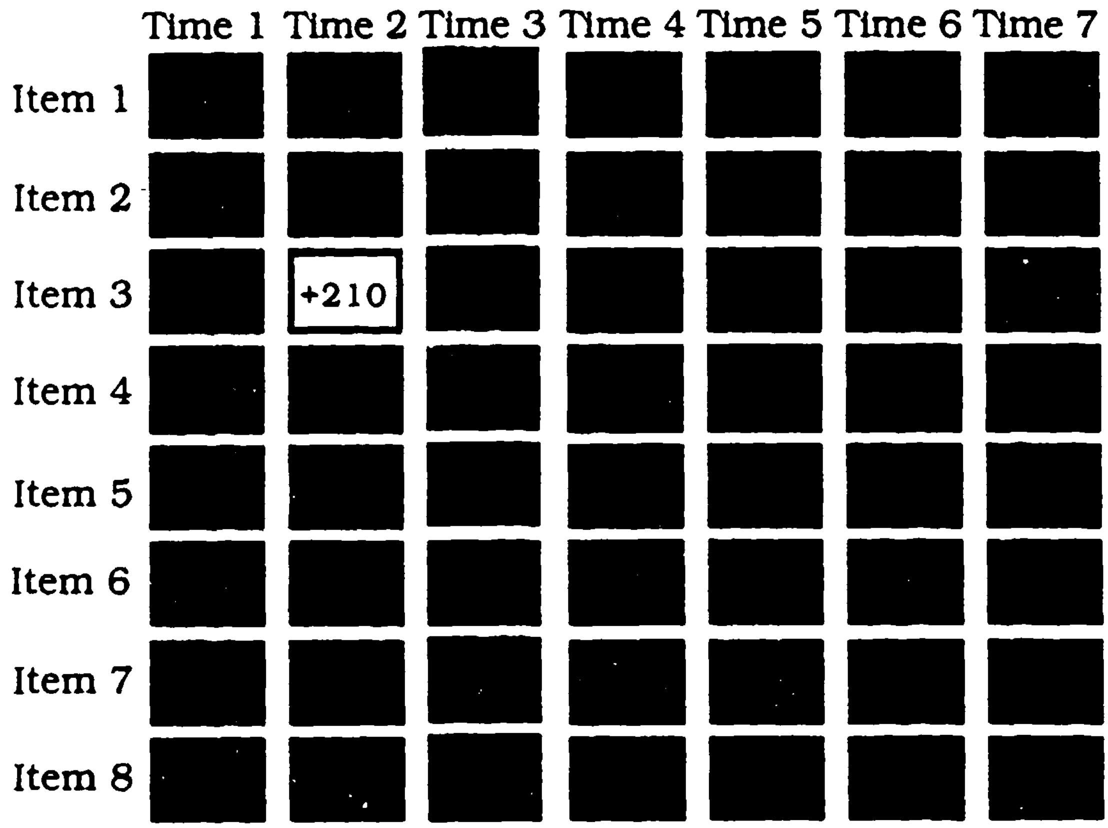

Figure 1. Sample Screen

NOVICE VERSUS EXPERT STRATEGIES

The first hypothesized strategy reflects differences in the quantity (QUANT) of information sampled. Expert decision-makers are efficient and, therefore, should search for data only where they expect to find significant information. The novice does not know what is most important and might reasonably attempt to gather as much information as possible. QUANT was thus defined by the total number of looks. It was predicted that novice decision-makers would make a larger total number of looks at the available data (higher QUANT) than would experienced decision-makers.

The second hypothesized strategy reflects predicted differences in memory among the groups. Chase and Simon[29] found that chess masters were better able to recall realistic chess positions than were beginning players. If, like chess masters, expert decision-makers are better able to remember realistic data than are novices, then experts should spend less time than novices reexamining data that have already been sampled (RELOOK). A measure

of the RELOOK strategy was defined as a count of the number of times the subject reexamines information previously seen. Novices were predicted to score higher in the RELOOK metric than experienced decision-makers. Note that the QUANT measure could be inflated by differences in memory for relevant data. If, as predicted, novices must reexamine previously inspected items, then this would add to the total number of looks. A post-trial questionnaire was used as a second measure of memory.

Novices can easily interpret only certain classes of information. For example, Larkin and Simon[30] found that experts included force vectors (abstract "objects") in their representation of physics problems while novices used only visible "objects." In the submarine, available raw data are analogous to force vectors. They represent measurements of abstract quantities that must be transformed before they can be interpreted as concrete indicators of spatial location. This transformation can be done by one of several computer algorithms. Considerable experience is required to interpret the raw data (as measured by sonar) in the submarine world and, therefore, instruction focuses on the more easily interpreted computer-processed data. It was observed, however, that experienced personnel make considerable use of the raw data, both to verify the reliability of the computed data, and to evaluate variables not apparent in the computed data, such as "crispness." In the present experiment, two classes of data were available: computer-processed and raw information. Therefore, it was predicted that experienced subjects would use the raw data class (CLASS) more frequently than the less experienced subjects.

In choice decision-making, the available information consists of stable values for various attributes (e.g., cost, location, effort) for each choice. For risky decision-making, the payoff or cost and probability of occurrence for each outcome is given a priori. In the dynamic, ambiguous situation, variables have values (e.g., direction, speed) that change over time. In some cases, a single look at most information sources is all that time will allow. In other scenarios, such as the submarine driver attempting to determine how to respond to a passive sonar target, there is time to sample each variable several times. In such situations, information on the rate and direction of change can help one predict the future state. However, this is a search strategy learned with experience. The frequency of examining the data history (that is, looking at the same variable, first at time 1 and immediately thereafter at time 2) was defined as a measure of the history (HIST) strategy (figure 1). It was predicted that experienced decision-makers would have a higher frequency of such HIST transitions than would inexperienced decision-makers.

In contrast, for the novice decision-maker, an efficient data-gathering strategy would be to examine data according to ease of access (EASE). If two items were physically adjacent, for example, transition between the two would be more efficient than between more distant variables (figure 1). A measure of the EASE strategy was defined as the frequency of examining physically adjacent information items, regardless of their logical connection to other strategies. On this measure, it was predicted that novice decision-makers would have higher frequency scores than experienced decision-makers. A second, corrected (EASE)-(HIST) measure of the EASE strategy also was defined because history information was adjacent in the experimental display and would otherwise be included in the EASE measure.

Research on novice-expert differences in chess[29] and physics[25] has shown that experts store information in well-defined, organized chunks. Such chunks are analogous to sets of related variables or values. Newell[12] has shown that sets of related items can be identified by a relatively short transition time (latency) between items within the set as compared with a longer transition time between sets. Two types of data sets can be defined. The first of these consists of a cluster of variables that provides related parts of a single chunk of information, that is, a part of the pattern of information stored in memory. Thus, one might expect experienced decision-makers to examine larger and more consistent sets of data (SET). This strategy was defined by a short, mean transition time between items and a high probability of transitioning among members of the set. One specific substrategy of SET is the case in which two (or more) items can provide a reliability check or confirmation for one another (RELI). An example of the use of this strategy would be taking data on the same variable from two or more sources, such as listening to two weather reports and then looking out the window. The RELI substrategy was defined as a look at a cluster of items whose members offer confirming data on the mutual reliability of the values, such as those offered by different sources of the same information. In the submarine, certain items of raw data can be used to judge the accuracy of the picture depicted by the computer-processed information. It was predicted that experienced decision-makers would have more easily identified SET and RELI chunks than inexperienced decision-makers.

INSTRUCTOR STRATEGIES

In the previous section, observable differences in information search strategies (and one substrategy) between novice and experienced decision-makers have been hypothesized. Much of the research into novice-expert differences has used problems from introductory texts that are very easy for the expert or for the proficient instructor. (See Larkin[25] for an exception.) On such problems, one could expect that the performance of a rule-bound subject would closely match that of an expert. Expertise, however, is best defined by effectiveness in situations that do not conform to easily enunciated rules. Only such situations should be expected to separate the rule-driven subject from the true expert with both rule-driven and experience-driven knowledge. To test the prediction that experts would differ from instructors primarily in difficult situations, two separate types of situations (one conforming to expected data patterns and one not conforming, i.e., more difficult) were presented within a single scenario. It was predicted that there would be greater differences in some of the information search strategies and, by implication, in decision-making performance in the portion of the scenario that did not conform to expected patterns of information.

Of particular interest in the difficult situation was the use of different classes of data. As noted earlier, raw data values can be used to evaluate the accuracy of computer-processed values. Thus, examination of the raw data can help the decision-maker detect anomalies within the data that may lead to false conclusions about the situation and, therefore, to poor decisions. The use of raw data (CLASS), however, is not the focus of instruction, and this is not featured prominently in the well-learned rule set

of submarine instructors. When faced with incongruous raw data, the instructor is likely to focus on the information that is best understood, namely, computed data. Therefore, it was predicted that instructors would make less use of the raw data class than experienced decision-makers, particularly in the difficult portion of the experiment.

Several lesser differences between instructors and the other two groups were predicted, across both easy and difficult portions of the scenario. First, in the submarine example, the use of confirming data (RELI) from several sources generally includes comparison of raw and computer-processed data and is a strategy largely learned from experience. Thus, it was predicted that instructors would not show interpretable and consistent RELI sets.

For strictly rule-driven decision-makers, some members of the SET clusters may be expected to differ from those of experienced people who have modified the strictly rule-driven expectancy sets to include or exclude some items, depending on their experience. Therefore, it was predicted that instructors would have some SET search patterns, but that these would differ from those of experienced submarine officers.

Use of the EASE and HIST strategies should depend on two factors: (1) the importance of history in the known rule set, and (2) where the individual falls in the novice-expert continuum. As history is an important feature of naval doctrine and instructors are very practiced with known rules, they should resemble experienced decision-makers in their use of the EASE and HIST strategies. There is no a priori reason to expect instructors to fall at one end of the RELOOK measure or the other. Therefore, it was predicted that they would fall between the novice and experienced subject groups. Lastly, as the rule-bound decision-maker has a set of expected sources of information, it was predicted that instructors would be similar to experts on their use of the QUANT strategy. A summary of predicted strategy use is given in table 1.

EXPERIMENTAL APPROACH

To test predicted differences in the use of these decision strategies, a process-tracing approach[3] was implemented. The information provided to the subjects consisted of selected portions of data from a submarine approach and attack exercise. These were arranged in a computerized matrix of time versus type of information and could be accessed by the subject with a mouse.

Information search is one stage in the decision process. The hypothesis that differences in search strategy would be related to level of experience was based on the assumption that the level of experience is related to the ability to integrate relevant information and make competent decisions. This assumption was examined by assessment of think-aloud verbal protocols and post-trial questionnaires. Verbal protocols can provide information on short-term and long-term goals, as well as on search strategy. Written questions were used to elicit data on post-trial information recall, ability to translate information from numeric to spatial format (a typical procedure for submariners), goals, and decision quality.

Table 1. Predicted Strategy Usage by Group

Strategy	Subject Groups		
	COs	Instructors	Students
QUANT	Low	Low	High
CLASS	Even	Mostly computed	Mostly computed
RELOOK	Low	Mid	High
HIST	High	High	Low
EASE	Low	Low	High
(EASE)-(HIST)	Low	Low	High
SET	Identifiable SET use	Some identifiable SETs	No SETs
RELI	Identifiable SET use	Very little	No SETs

METHOD

The method used for testing is discussed in the following sections.

Subjects

The subjects were 12 current or retired Navy submarine officers. They were divided into three groups of four subjects each, with each group composed of a former submarine CO, an instructor in various aspects of submarine operations, and two students in the first course for submarine officers. Each of the COs had a minimum of 3 years experience in command of the submarine and 5 years at sea. The instructors had a minimum of 7 months as instructors and 3 years of sea experience in the role of junior officer. Students had not yet served aboard a submarine. While this is not a large number of subjects, the sample does represent an unusually large percentage of the total submariner population.

Apparatus

The apparatus included a Zenith 181 portable computer with a PC Mousesystems mouse, a microrecorder, and a questionnaire with judgment questions for each of the three legs of data and additional end-of-session questions. (A leg is a period during which the ship moves at a relatively constant course and speed while data are gathered on a sonar contact. Because of the under-constrained nature of passive sonar data and the uncertain nature

of sound paths underwater,[32] several legs are usually necessary to determine the location and motion of the source of a passive sonar contact.) To test the effects of the consistent, predictable information patterns versus inconsistent, unpredictable patterns, the third leg contained an anomaly that served to make interpretation difficult, within the context of the usual rule set.

The matrix schema of MOUSELAB[33] was used to present data to the subjects. In the matrix schema, information is presented in closed boxes identified by row and column labels (figure 1). The information is visible to the subject only when the mouse cursor is moved into the box. When the cursor leaves the box, the contents are again hidden. The program collects and stores the items examined and the entry and exit time for each item.

Eight items of information typically available in a submarine command center could be accessed by moving the mouse into the desired box. These items were divided into two classes.

The computer-processed (system solution) data were as follows:

1.	Computed bearing	(CBy)
2.	Computed bearing rate	(CDBy)
3.	Computed course and speed	(CC/S)
4.	Computed range to target	(CR)

The raw sonar data were as follows:

5.	Raw bearing	(RBy)
6.	Bearing rate from RBy	(RDBy)
7.	Signal-to-noise ratio	(SNR)
8.	Sonar tracker identification	(Tr).

To evaluate the use of the RELI substrategy, sets {1,5} and {2,6} contained redundant information in the computed and raw forms. Furthermore, the accuracy of item 4 could be checked against either or both members of the set {2,6} because bearing rate is a rough indicator of target range. The data were abstracted from an at-sea fleet exercise, and subjects were told that they might contain human and/or machine errors or ambiguities. Using exercise data gave the experiment face validity for the subject population and avoided the expectation of a perfectly knowable situation.

At the end of each segment (leg) of the exercise, a written questionnaire was used to gather data on subject recall of situational information and decisions. These questions included recall of specific data, estimates of data accuracy, translation of numeric data into a typical submarine graphic representation, orders for any actions, and questions about goals.

Procedure

Subjects were briefed about the purpose of the experiment and source of the data. Instructions were presented on the screen, including definition of

labels and organization of the scenarios. A sample screen was presented with nonsense data to allow the subjects to practice using the mouse. A second practice scenario was displayed with fictional data for further practice with the experimental paradigm. If necessary, the experimenter answered any questions before the three experimental trials began. Subjects were requested to "think aloud" throughout the experiment. They were reminded of this request during the practices, but they were not reminded during the actual experimental trials. The three trials represented three legs of a submarine search mission in which the submarine was attempting to determine the location of an unknown sonar target. Subjects were given 90 seconds to examine the data available from each leg. At the completion of each leg, the subjects were instructed to answer questions in the response booklets. After answering the questions, they were told what maneuver the submarine had made during the exercise and were able to examine data for the next leg. At the completion of the third leg, additional questions assessed the subjects' overall evaluation of the situation and experience level.

RESULTS

The MOUSELAB output files were converted for analysis as (1) dwell time in each box (duration), (2) transition time between boxes (latency), and (3) item selection frequency. Durations of 0.05 second or less were eliminated as these occurred too fast for the box to open, and they were assumed to be the result of the mouse passing through a box rather than the subject examining data in that box. Results, analyzed for differences among subject levels for each strategy, are given in the following sections.

QUANT, RELOOK, HIST, EASE and (EASE)-(HIST)

Frequencies for each of these groups are given in table 2. All data were analyzed using the QUANT proportions to calculate expected frequencies. (See figure 2.)

Table 2. Frequency of Strategy Usage by Group

	Subject Group		
Strategy	COs	Instructors	Students
QUANT	959	1060	1282
RELOOK	562	616	817
HIST	774	680	809
EASE	937	970	1196
(EASE)-(HIST)	163	290	391

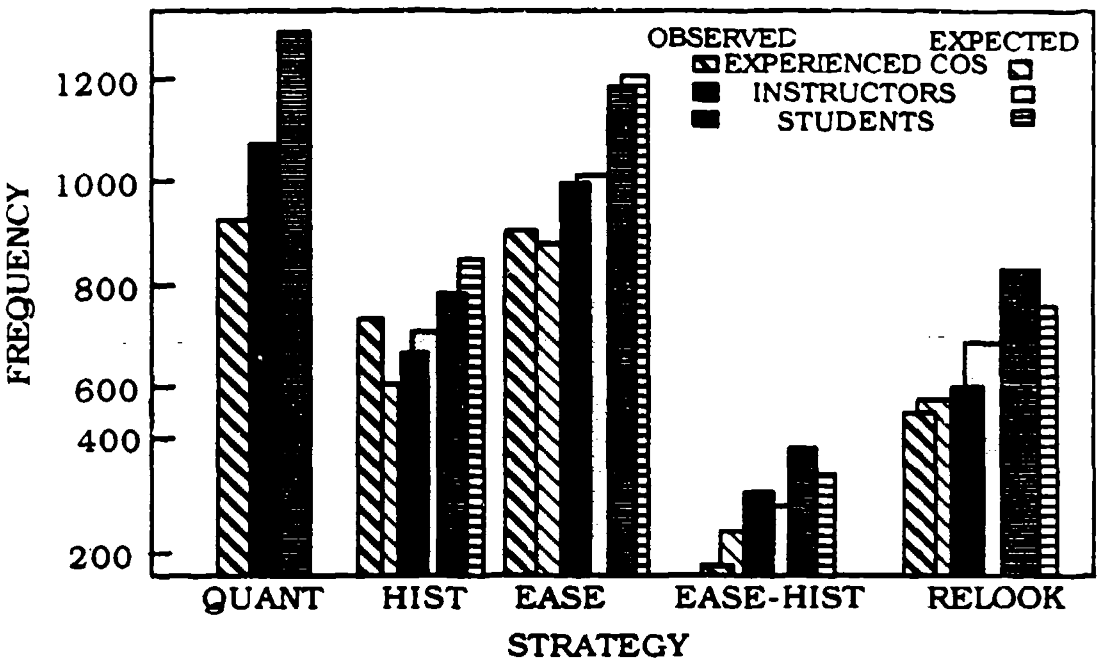

Figure 2. Frequency of Looks for QUANT, HIST, EASE, (EASE)-(HIST), and RELOOK Strategies by Group

Significant differences were found among groups on all measures except for RELOOK and EASE (uncorrected for HIST): $X^2(4) = 16.647$, probability (p) < 0.001 for QUANT; $X^2(2) = 3.679$, not significant (n.s.) for RELOOK; $X^2(2) = 30.086$, $p < 0.001$ for HIST; $X^2(2) = 2.193$, n.s. for EASE; and $X^2(2) = 40.916$, $p < 0.001$ for (EASE)-(HIST). Notice especially that students made significantly more looks at data (greater QUANT) and made greater use the (EASE)-(HIST) strategy than did the COs. In contrast, although they used the fewest looks (lowest QUANT), COs made greater use of the HIST strategy than either of the other groups. Instructors fell between the two other groups, resembling COs in their use of the QUANT strategy and the students in their relative use of the (EASE)-(HIST) strategy.

These findings support the prediction that the three groups would differ in their use of information-gathering strategies. These differences were not just quantitative, however. They were also apparent in the verbal protocols and written responses. Such findings imply differences in the decision process itself and lend credence to the assumption that (in this case) good, efficient decision-making is, first, good situation assessment.

Leg and CLASS

The previous analyses were across all legs. QUANT, EASE, (EAST)-(HIST), and RELOOK strategy usage was not predicted to differ by leg. Of these, only QUANT was found to differ significantly by leg (figure 3). Table 3 shows the leg and CLASS frequencies for the three groups. There were significant

Table 3. Frequency of Looks by Leg and CLASS by Group

CLASS	Subject Group COs	Instructors	Students
Leg 1			
Computed	158	239	331
Raw data	139	129	166
Leg 2			
Computed	123	245	238
Raw data	189	92	160
Leg 3			
Computed	183	253	232
Raw data	167	102	155

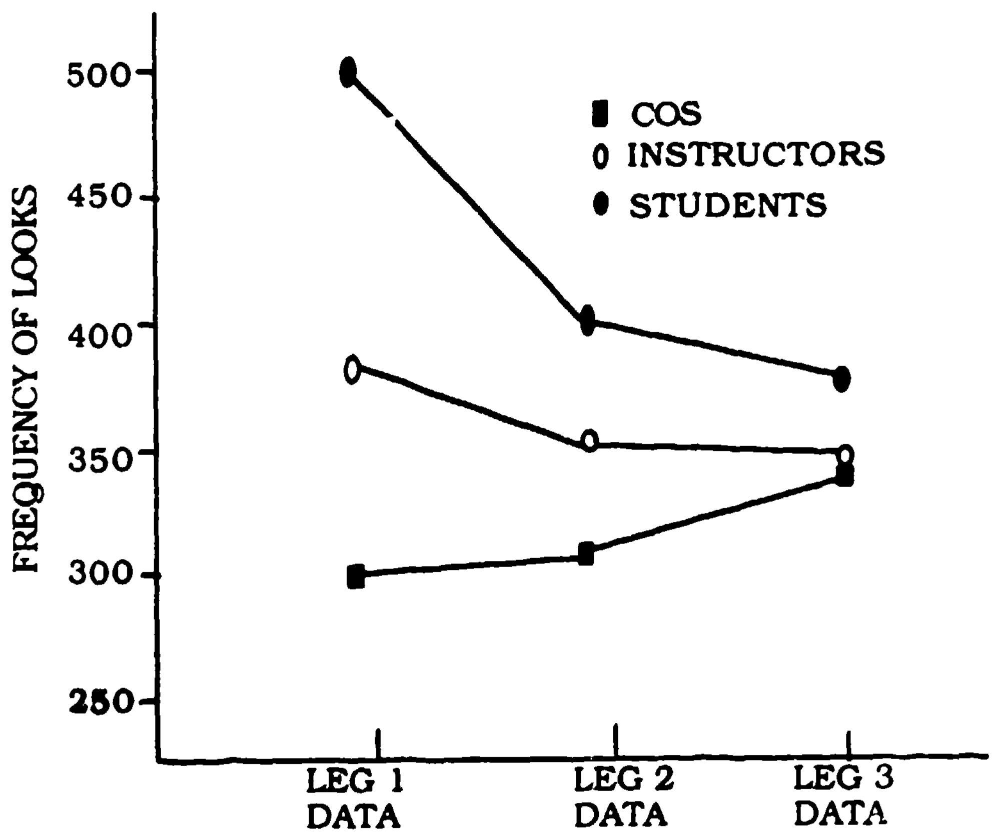

Figure 3. Frequency of Looks by Leg and Group

differences among the groups on both the frequency of looks by leg $\chi^2(4) = 16.647$, $p < 0.001$; and by CLASS $\chi^2(2) = 40.813$, $p < 0.001$. As can be seen by examining figures 4 and 5, the experienced officers were more even in their use of available data both across legs and across CLASSes of data. Of particular interest are the differences in search strategy on the more difficult leg 3. Note that instructors relied heavily on the computed data while experienced COs continued their even use of both classes of data (figure 5).

SET and RELI

The SET and RELI strategies were defined by groups of one or more short latency transitions, separated by longer transitions. Because of the small number of subjects, cluster analytic methods were not appropriate to examine these data; however, state transition diagrams (figures 6 through 8) for each of the three groups were drawn using mean transition latencies. These figures show these transition times between choices, as well as frequency of looks at each item for each subject group. The relatively structured pattern of information usage for the CO group (figure 6) as compared with the relatively unstructured pattern for the student group (figure 8) should be noted. Items 1 to 4 composed the computed data class and items 5 to 9, the raw data class. Although only the group patterns are given on figures 6 through 8, note that each individual subject had a different pattern of transition times. As described by Newell,[12] "Problem spaces imply that ranges of possible behaviors are to be expected... The same subject on repeated occasions will exhibit a range of behavior, even though working in the same space."

When figures 6 through 8 were shown to subjects during a post-experiment debriefing, they were able to explain all the COs' clusters and all but one of the instructors' clusters. They noted that the instructors' patterns of data usage showed a concentration on processed data, as compared with the COs' more balanced information usage. They did not find this surprising because instructors teach the operation of the processing system, not the use of data (raw and processed) to make decisions. They also noted the lack of identifiable clusters in figure 8, the students' transition diagrams.

Questionnaire and Verbal Protocol Data

The between-leg questionnaire data were analyzed for accuracy of recalling the given data and for accuracy in translating the numeric information into graphic representations. The mean accuracy of recall and translation to graphic format can be found in table 4. Note especially the high positive correlation between the two tasks for the CO group and the negative correlation for the other two groups. This may be an indication of the greater integration skill or more integrated knowledge sets that experience confers, over and above the use of rules without experience.

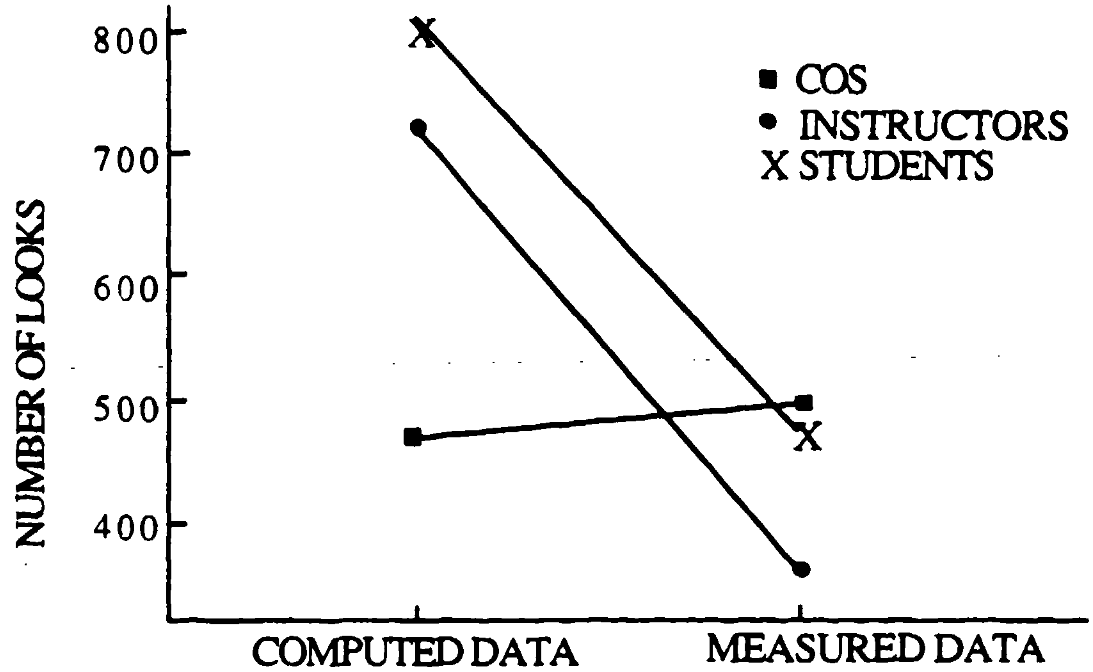

Figure 4. Frequency of Looks by CLASS and Group

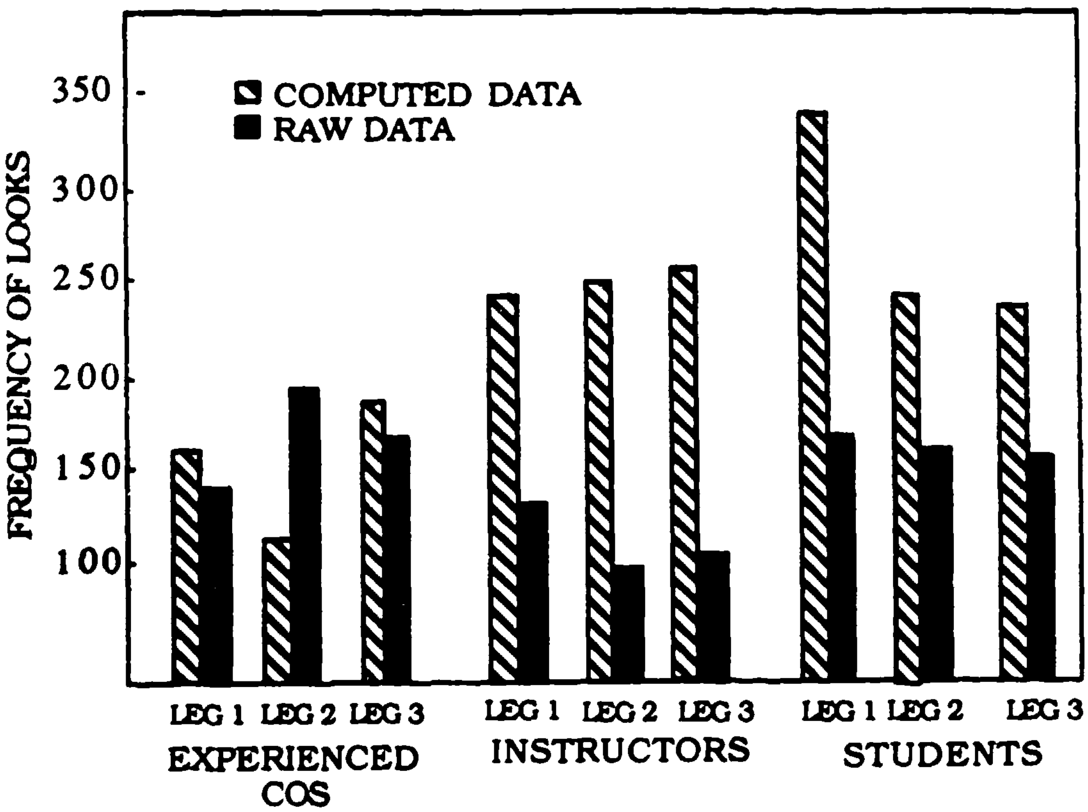

Figure 5. Frequency of Looks by CLASS and Leg for Each Group

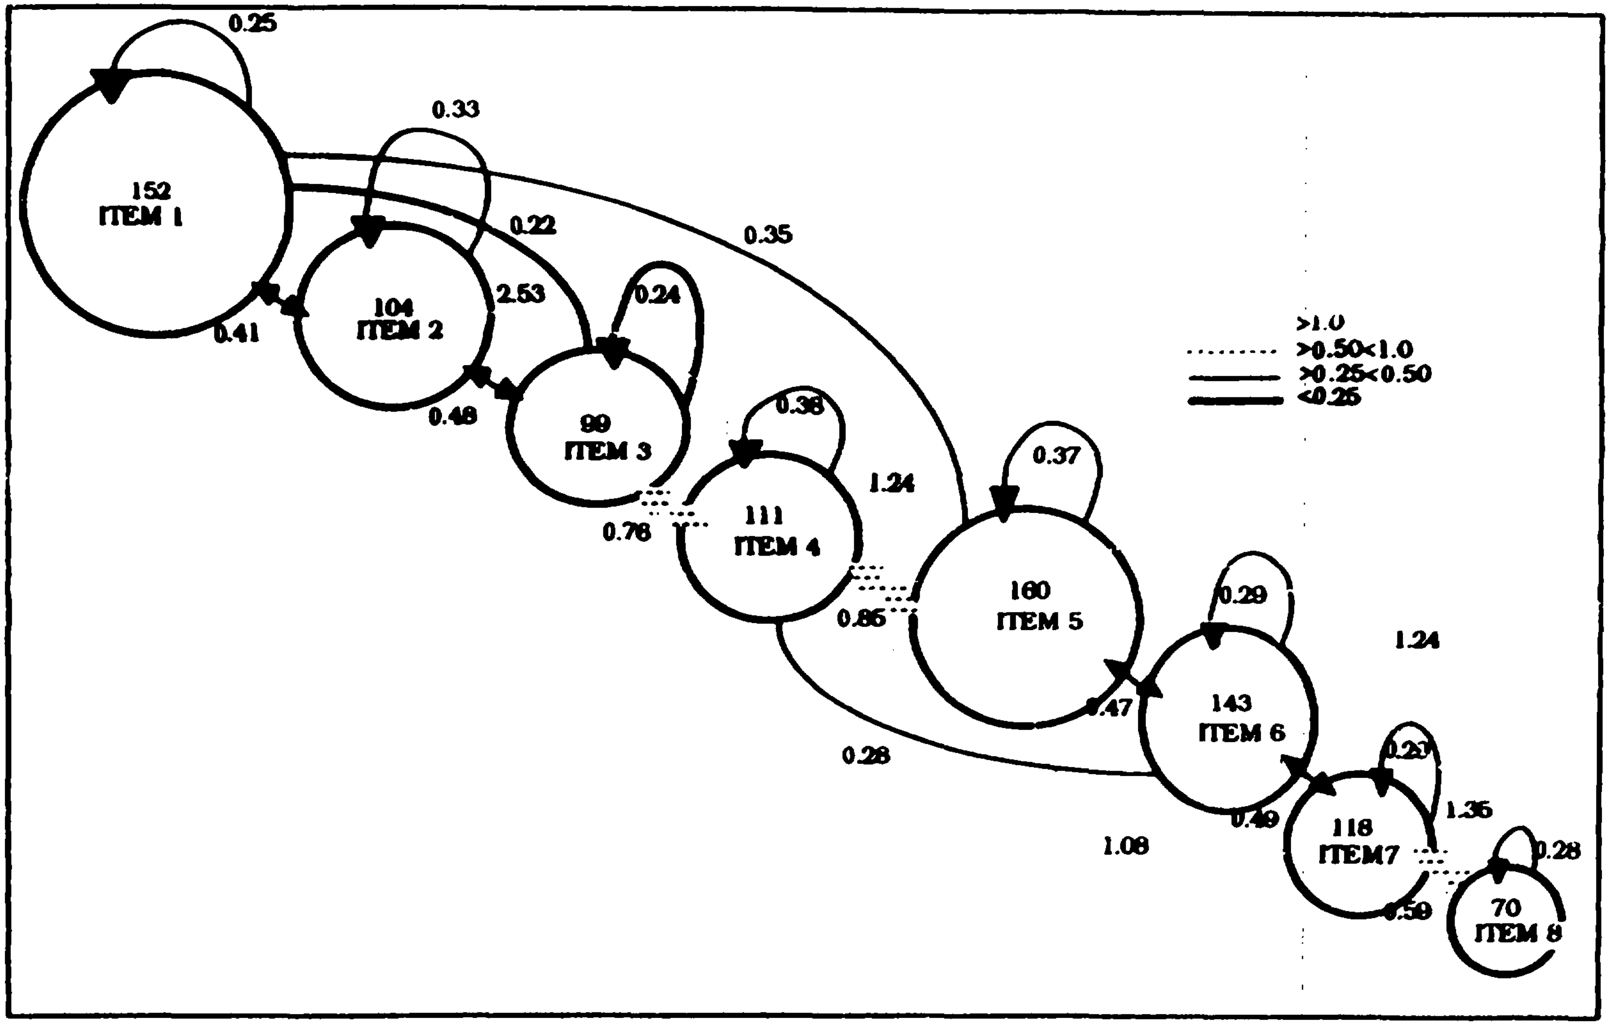

Figure 6. Frequency of Looks at Each Item and Mean Latency for Transition (in seconds) for Expert Subjects

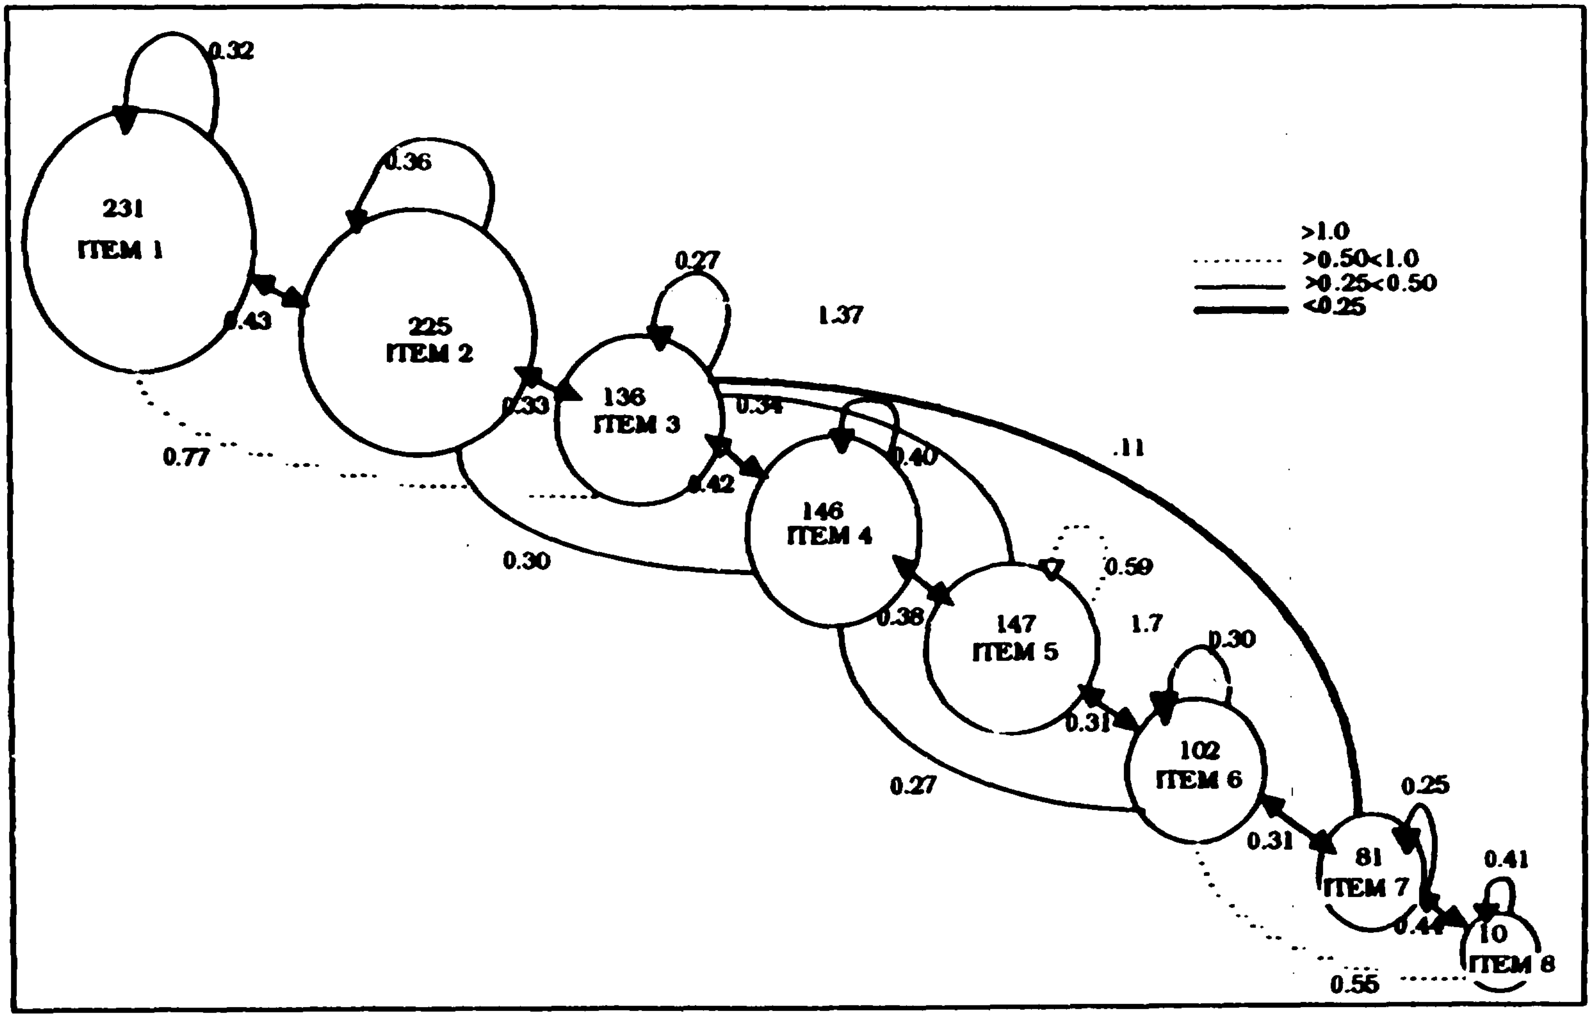

Figure 7. Frequency of Looks at Each Item and Mean Latency for Transition (in seconds) for Instructors

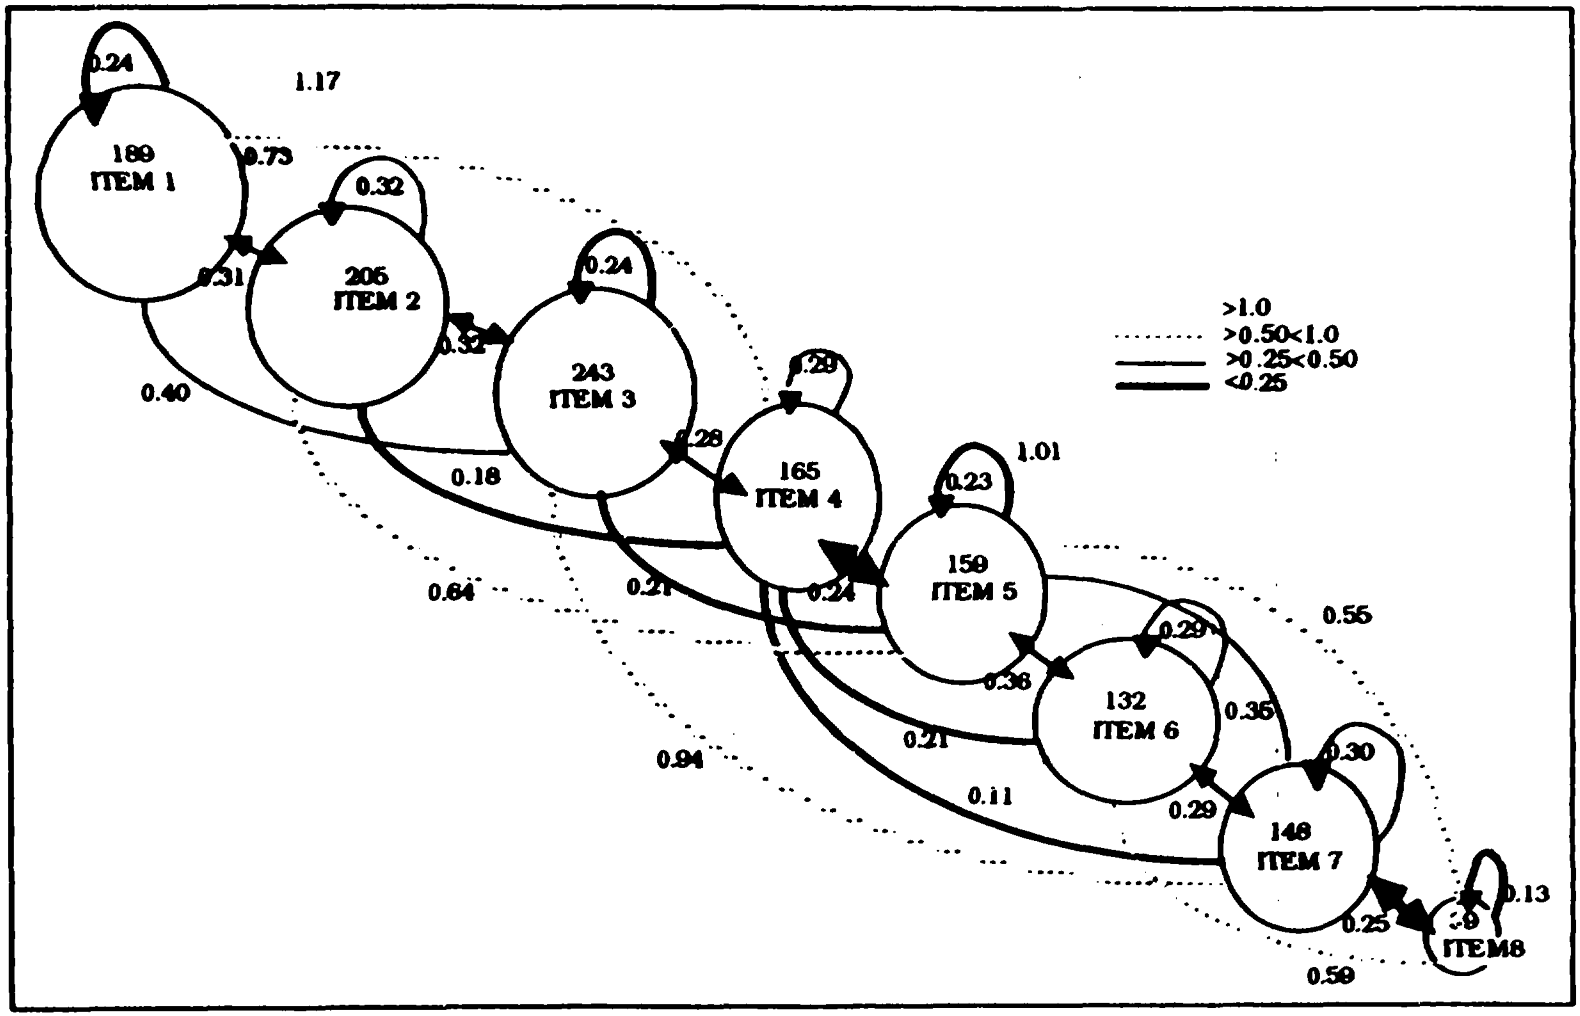

Figure 8. Frequency of Looks at Each Item and Mean Latency for Transition (in seconds) for Students

Table 4. Accuracy of Recall and Translation into Line-of-Sight Diagram

	COs		Instructors		Students	
	*M	SD	*M	SD	*M	SD
Graphic Translation	73%	19	62%	17	49%	22
Data Recall	62%	27	31%	14	61%	28
Correlation	+0.97		-0.79		-0.99	

*M = Mean
SD = Standard deviation

Questionnaire responses also were analyzed by an independent expert for quantity and quality of reasons given. Because of the small number of subjects and relatively large intergroup variability, no statistically significant conclusions could be drawn from these data. The COs gave an average of 8 reasons for their actions and averaged 3.25 out of 4 assessed quality points. Instructors averaged 7.3 reasons and 3.1 quality points. In contrast, the students averaged 4.8 reasons and 2.5 quality points.

Questionnaire responses also were used to assess decision responses made at the end of each leg. A team of experts in naval doctrine evaluated the given scenario information and determined the appropriate response(s) at the end of each leg. These were compared then with the decision responses of the subjects. The responses of the COs and instructors approached the recommended responses on the first two legs. Thus, rules were sufficient to make good responses for the predictable portion of the scenario. This judgment supports the original assumption that the effect of experience is to alter the application of standard rules. On the difficult third leg, the CO group appeared to respond more "correctly," according to expert opinion. This may be because they were more likely to examine both computed and raw data, especially in item 8. On the third leg, there was a mismatch between the two classes of data (processed and raw). Item 8 provided an indication of the reason for mismatch and an indication of the seriousness of the problem. Lastly, there was more apparent variation among members of the experienced CO group, both in terms of response actions and information search patterns, at this point. This is not surprising because training (i.e., learning the rules) is a fairly uniform process among submarine officers. Learning from experience, however, varies greatly among experienced officers because of the differences between duty assignments.

The think-aloud tapes of the experimental sessions were replayed and statements coded into the following four categories: reading and interpreting data, stating procedures, making evaluations, and making goal statements. Differences among the groups were apparent on these tapes. The inexperienced subjects spoke very little during the entire session and when they did, it was

frequently inaudible. They mainly read the data and did not include evaluative or goal statements. The other two groups used all four types of statements. Moreover, the experienced COs were more likely to give a hierarchy of goal statements, while instructors were more likely to name procedures or read/interpret data. During the third leg, the instructors fell noticeably quieter.

In summary, the questionnaire and verbal protocol data provided support for the behavioral data. All three data collection methods pointed to similar conclusions, and each provided some unique detail helpful for the interpretation of the others.

DISCUSSION OF SEARCH CHARACTERISTICS

Clear differences emerged among the groups in the use of information-gathering strategies, the quantity and quality of thinking aloud, recall of information without available visual cues, translation of numeric data into graphic representations, and the action decision itself. A summary of these differences can be found in table 5. Together, all available data suggest that the three classes of subjects do differ, not only in search strategy but also, by implication, in how they approach a decision problem. Such a pattern of differences is the first step in understanding how information search plays a part in decision-making and in how expertise influences information search. Recognizing this pattern of differences is vital because it allows the designer of an adaptable decision aid to recognize and differentiate user characteristics and adapt to the user differences. It is hoped that this adaptability will better support the decision process for all users. In the following sections, the information search characteristics of the three groups are discussed in detail.

STUDENTS

The differences between the two extremes (expert and novice decision-makers) on all measures were, not surprisingly, the most clear findings of this experiment. Students employed a search strategy based on gathering as much information as possible. To this end, they had a very high proportion of transitions between adjacent items. They also frequently returned to previously examined information. The payoff for this strategy was their ability to accurately recall information at the end of each leg. This finding is in contrast to Chase and Simon,[29] who found that novices were significantly poorer at recall than chess masters. Note, however, that the number of values to be recalled in the present experiment was only 4 for each leg while the number of chess positions varied from 14 to 24.

This quantity-oriented strategy may have impacted the novice subjects' ability to integrate the available information. Such an integration would likely take more time for the novice than the more knowledgeable person, but the novices actually allowed somewhat less time because they were so busy moving from item to item. The lack of verbal protocol data is another indication of the novices' emphasis on gathering the greatest quantity of

Table 5. Actual Strategy Usage by Group

Strategy	Subject Group		
	COs	Instructors	Students
QUANT	Low numbers	Moderate	High numbers
CLASS	Even use	More computed than raw	Far more computed than raw
RELOOK	Moderate	Moderate	Higher
HIST	High	Even	Low
EASE	Moderate	Moderate	Higher
(EASE)-(HIST)	Low	Higher	Highest
SET	{1,2,3} {5,6,7}	{3,5}	Unclear {3,7}
RELI	{4,6} {1,5}	{2,4,6}	Unclear

information at the expense of everything else. For the novice, the effort to think aloud may have been distracting or may have required more cognitive effort than was available, since they were already fully engaged in gathering situational information. In contrast, both instructors and experienced COs are accustomed to explaining what they are doing as a teaching device and this may account for the differences in quantity of verbal data.

Paradoxically, although novices looked at the greatest quantity of data, they showed a limiting and disproportional concentration on computed data. As was previously noted, the computer-processed data are the focus of most course work and are easier to understand than the raw data. Thus, even for the novice trying to ingest as much information as possible, there was an apparent selection process. It was not the same as that of the expert, but it did guide information gathering toward that information the subject expected to find most useful. Such an information search strategy implies an expected pattern of usefulness in situational information. This pattern of usefulness could have developed from direct positive instruction or by observation of the model provided by instructors. In the present experiment, the search patterns of the novice can be compared and constrasted with those of his instructors.

On the first two legs, instructors resemble experts, not students. Instruction, however, generally focuses on the aspects of the task the instructors consider most important, most fundamental, or easiest to grasp. Teachers do not teach everything they know; they simplify. It cannot be judged whether teaching a more expert-like, information-gathering strategy

would move students up the experience ladder more quickly. Alternatively, perhaps presenting information in a more expert-like mode will improve the performance of novice decision-makers.

INSTRUCTORS

The rule-bound instructor group appeared to have the least within-group variation on most measures, since well-formulated rules are intended to limit variability. The goal of the rules written into naval doctrine is to anticipate most likely situations and provide response guidance. Experience builds on the standard rules by providing a basis for responding to nonstandard situations. Thus, it was predicted that instructors would generally resemble experts on the standard legs (1 and 2) but differ on the difficult leg (3). This prediction was partially supported. Overall, instructors were similar to COs in the use of the QUANT and SET strategies, and similar to students in their use of the HIST, EASE, (EASE)-(HIST), and CLASS strategies. Moreover, their use of the QUANT strategy changed between legs 1 and 2. Interestingly, instructors were the least likely of all groups to examine the key item on leg 3 (tracker) and, therefore, to recognize the cause of the anomaly. Only one member of this group looked at the critical item 8 during that leg, and then only once. In comparison, novices looked at item 8 a total of 5 times and experts a total of 17 times. Recall that the decision performance of instructors was judged most similar to the naval doctrine of the groups for legs 1 and 2. It fell below that of experts only for the more difficult leg 3. Lastly, although accustomed to explaining as they perform, the instructors were noticeably quieter on the third leg. The contrast in patterns of responses between the legs supports the assumption that information-gathering behavior is critical to decision-making. If instructors had used an expert-like information-gathering strategy on the third leg, would they have shown a better understanding of the situation and made better decision responses? Could an adaptable decision support system help bridge this gap? The building of Ranger may help answer this and other similar questions.

EXPERIENCED COMMANDING OFFICERS

Experts were both the most even and most clearly structured group in their examination of data. Their basic strategy focused on selectively examining information. Evenness, history, and critical sets of related data figured prominently in their search patterns. They made significantly fewer looks at the available information but were good at recalling and transforming data at the end of each leg. They also spent more time talking about the process and goals than did the other groups. This picture of experts as systematic and efficient gatherers of information concurs with other studies that show experts as recalling larger chunks of information than novices[29] perceiving both surface and deep problem-structure,[24] representing problems better than novices,[25,27] and solving problems faster and with different heuristics than novices.[34] Clearly enunciated goals and subgoals were reported on both verbal protocols and questionnaire responses. After each leg, experts were able to recall critical data and transform numeric values into spatial representations. This spatial transformation is intuitively easier to interpret and therefore to guide maneuver decisions in the submarine

world than is the untransformed sonar information. Lastly, the critical third leg showed experts as prepared to deal with difficult and ambiguous situations without altering their basic information search strategy. By not retreating to a less efficient strategy, they were able to pinpoint the anomaly and compensate for it, which is the essence of expertise.

SUMMARY AND RECOMMENDATIONS

Some differences in decision-making strategies among submariners due to differences in levels of experience have been explored in this report. While these differences have been expressed in terms of information-gathering strategies, some of which are specific to the submarine decision problem, they also can be generalized to other dynamic, ambiguous decision situations. With additional specification of these differences, more flexible decision support may be provided for all types of decision-makers. It is yet to be examined whether the structured information search strategies of the instructor would help the novice and if the balanced search strategies of the expert would help the rule-bound instructor become more flexible when dealing with difficult situations.

These results have shown several ways in which an expert system could sense and adapt to user differences in level of experience, and also have shown some requirements for such a system. The first of these requirements is that the system have information requests to use as clues about the user. That is, the user must interact with the system, collecting situational information for cooperative decision-making, rather than just request an expert system solution. Secondly, the system must have knowledge about the reasonable data sets or information request paths for the problem under consideration. This means that it must have a knowledge about possible situations. It must also offer the user the opportunity to examine both computed solutions and the raw data that contribute to the solutions as this has been shown to match the information-gathering strategy of experienced decision-makers. Furthermore, it must be able to store data about user actions and reason across these, just as it does situational data. Lastly, it must be able to adapt its output to the needs of the current user.

How to do this last recommendation is only suggested by this study. Sufficient data have been gathered, however, to define likely user strategies and suggest possible adaptations. For example, the user who, after some period of time, does not request items from known sets, should be given the unrequested items with the requested ones. Moreover, since such a user is likely to be an inexperienced decision-maker, the relationships between requested and presented items should be made explicit. In addition, when anomalies occur and the user does not request the relevant data within a reasonable time, he should be alerted to the problem. The ways in which these recommendations have been implemented in Ranger and in which they are still to be improved are reported in Cary.[2]

As this was an interdisciplinary collaborative project, there are also recommendations for extensions of the psychological research. Much has been learned about differences in information-gathering strategies for decision-

making, but there have been as many questions raised as answered. Continued research is needed to investigate how specific environmental variables affect decision-maker behavior. Research into the differences between decision-makers with average experience and those who are superior at some specific subset of the task also could contribute to knowledge of both how to develop aids and how to train for superior performance. In addition, when Ranger is extended to comprise all the intended functionality, the premise of this work, that adaptability can help improve human performance and system acceptance, should be formally evaluated. The concept of adaptability, which contributes so much to human-human communications, has yet to be tested in human-computer communications. Ranger offers a ready facility for such a test.

REFERENCES

1. E.A. Isaacs and H.H. Clark, "References in Conversation Between Experts and Novices," Journal of Experimental Psychology: General, vol 11, 1987, pp. 26-37.

2. J. Cary, "Adaptable Expert Systems for Submarine Command-Level Decision Aiding," NUSC Technical Report, Naval Underwater Systems Center, Newport, RI (in preparation).

3. J.W. Payne, "Contingent Decision Behavior," Psychological Bulletin, vol 92, 1982, pp. 382-402.

4. J. Klayman, "Simulations of Six Decision Strategies: Comparisons of Search Patterns, Processing Characteristics, and Response to Task Complexity," unpublished manuscript, University of Chicago, Graduate School of Business, Center for Decision Research, Chicago, IL, 1982.

5. J. Klayman, "Analysis of Predecisional Information Search Patterns," in Analyzing and Aiding Decision Processes, P. Humphreys, O. Svenson, and A. Vari (eds.), North-Holland, Amsterdam, 1983.

6. H.J. Einhorn and R.M. Hogarth, "Behavioral Decision Theory: Processes of Judgment and Choice," Annual Review of Psychology, vol 32, 1981, pp. 53-88.

7. H.J. Einhorn and R.M. Hogarth, "A Contrast/Surprise Model for Updating Beliefs," Report No. 13, Center for Decision Research, Chicago, IL, 1985.

8. H.J. Einhorn and R.M. Hogarth, "Judging Probable Cause," Psychological Bulletin, vol 99, 1986, pp. 1-17.

9. H.J. Einhorn, D.N. Kleinmuntz, and B. Kleinmuntz, "Linear Regression and Process-Tracing Models of Judgment," Psychological Review, vol 86, 1980, pp. 465-485.

10. O. Svenson, "Process Descriptions of Decision Making," Organizational Behavioral and Human Performance, vol 23, 1979, pp. 88-112.

11. A.D. deGroot, Thought and Choice in Chess, Mouton, The Hague, Netherlands, 1965.

12. A. Newell, "Reasoning, Problem Solving, and Decision Processes: The Problem Space as a Fundamental Category," in Attention and Performance VIII, R. Nickerson (ed.), Erlbaum, Hillside, NJ, 1980.

13. Y. Anzia and H.A. Simon, "The Theory of Learning by Doing," Psychological Review, vol 86, 1979, pp. 124-140.

14. M.T.H. Chi, P.J. Feltovich, and R. Glaser, "Categorization and Representation of Physics Knowledge by Experts and Novices," Cognitive Science, vol 5, 1981, pp. 121-152.

15. W. Edwards, "The Theory of Decision-Making," Psychological Bulletin, vol 12, 1954, pp. 473-498.

16. W. Edwards, "Dynamic Decision Theory and Probabilistic Information Processing," Human Factors, vol 4, 1967, pp. 59-73.

17. D. Kahneman and A. Tversky, "Choices, Values, and Frames," American Psychologist, vol 39, 1984, pp. 341-350.

18. S.S. Kirschenbaum, "A Multistage, Multilevel Conceptual Model of Human Decision-Making," NUSC Technical Memorandum No. 86-2083, Naval Underwater Systems Center, Newport, RI, 30 September 1986.

19. J.G. Wohl, E.E. Entin, and J.S. Eterno, Modeling Human Decision Processes in Command and Control, Report. No. TR-137, Alphatech, Inc., Burlington, MA, 1983.

20. D. Kahneman and A. Tversky, "Subjective Probability: A Judgment of Representativeness," Cognitive Psychology, vol 3, 1972, pp. 430-454.

21. A. Tversky and D. Kahneman, "Belief in the Law of Small Numbers," Psychological Bulletin, vol 76, 1971, pp. 105-110.

22. A. Tversky and D. Kahneman, "Availability: A Heuristic of Judging Frequency and Probability," Cognitive Psychology, vol 5, 1973, pp. 207-232.

23. J. Larkin, J. McDermot, D.P. Simon, and H.A. Simon, "Expert and Novice Performance in Solving Physics Problems," Science, vol 208, 1980, pp. 1335-1345.

24. A.H. Schoenfeld and D.J. Herrmann, "Problem Perception and Knowledge Structure in Expert and Novice Mathematical Problem Solvers," Journal of Experimental Psychology: Learning, Memory, and Cognition, vol 8, 1982, pp. 484-494.

25. J.H. Larkin, "The Role of Problem Representation in Physics," in Mental Models, A. Gentner and A.L. Stevens (eds.), Erlbaum, Hillsdale, NJ, 1983.

26. W.B. Rouse and N.M. Morris, "On Looking into the Black Box: Prospects and Limits in the Search for Mental Models," Psychological Bulletin, vol 100, 1986, pp. 349-363.

27. J.H. Larkin, "Understanding, Problem Representations, and Skill in Physics," in Thinking and Learning Skills: Research and Open Questions, vol 2, S.F. Chipman, J.W. Segal, and R. Glaser (eds.), Erlbaum, Hillsdale, NJ, 1985.

28. D. Nobel, D. Boehm-Davis, and C. Grosz, "Schema-Based Model of Information Processing for Situation Assessment," Contract No. N00014-84-C-0484, Engineering Research Associates, 8618 Westwood Center Drive, Vienna, VA, 1986.

29. W.G. Chase and H.A. Simon, "Perception in Chess," Cognitive Psychology, vol 4, pp. 55-81, 1973.

30. J.H. Larkin and H.A. Simon, "Why A Picture is Worth Ten Thousand Words," Department of Psychology, Carnegie-Mellon University, Pittsburgh, PA, 1985.

31. I. Vessey, "Expertise in Debugging Computer Programs: An Analysis of the Content of Verbal Protocols," IEEE Transactions on Systems, Man, and Cybernetics, vol SMC-16, 1986, pp. 621-637.

32. R.J. Urick, Principles of Underwater Sound (3rd Edition), McGraw-Hill, NY, 1983.

33. E.J. Johnson, J.W. Payne, D.A. Schkade, and J.R. Bettman, "Monitoring Information Processing and Decisions: The MouseLab System," Carnegie-Mellon University, Graduate School of Industrial Administration, Pittsburgh, PA, 1986.

34. R. Bhaskar and H.A. Simon, "Problem Solving in Semantically Rich Domains: An Example for Engineering Thermodynamics," Cognitive Science, vol 1, 1977, pp. 193-215.

INITIAL DISTRIBUTION LIST

Addressee	No. of Copies
COMSUBDEVRON 12 (CAPT Almon)	1
DARPA (C. Stewart)	1
CNR (OCNR-1142 (S. Chipman, H. Hawkins, J. O'Hare, W. Vaughan), OCNR-1211 (J. Smith), OCNR-20 (S. Collyer), OCNR-23 (A.Faulstich, D. Houser))	8
NOSC (Code 440 (J. Grossman))	1
SUBSCOL (Code 812, Code 10)	2
DTIC	12
CNA	1